



United States Patent Application

**METHOD AND APPARATUS FOR SHORTENING DISPLAY LIST
INSTRUCTIONS**

Inventor: GORAN DEVIC
10401 Doc Holliday Trail
Austin, TX 78753

I hereby certify that this document is being deposited with the United States Postal Service as Express Mail in an envelope
marked as "Express Mail Post Office to Addressee" Mailing
Label Number PM 28326240045
addressed to the Commissioner of Patents and Trademarks,
Washington, D.C. 20231 on this date March 14, 1997

Steven A. Shaw
Typed or printed name of person mailing this document

[Signature] 3/14/97
Signature Date

Assignee: **Cirrus Logic, Inc.**
3100 West Warren Ave.
MS 521; Legal/IP
Fremont, CA 94538
(510) 252-6242

08818053-031497

Field of Invention

The present invention relates generally to a graphics system for personal computers. More particularly, the present invention relates to a method and apparatus for shortening display list instructions in a graphics processor.

Description of the Related Art

Sophisticated graphics packages have been used for some time in expensive computer design and graphics systems. Increased capabilities of graphics controllers and display systems, combined with standardized graphics languages, have made complex graphics functions available in even the most routine applications. For example; word processor, spread sheets and desktop publishing packages now include relatively sophisticated graphics capabilities. Three-dimensional (3D) displays have become common in games, animation, and multimedia communication and drawing packages.

The availability of sophisticated graphics in PCs has driven a demand for even greater graphics capabilities. To obtain these capabilities, graphics systems must be capable of performing more sophisticated functions in less time to process greater amounts of graphical data required by modern software applications. In particular, there is a continuing need for improvements in software algorithms and hardware implementations to draw three-dimensional objects using full color, texture mapping and transparency blending.

Improvements have been made in the hardware realm. Graphics processors and accelerators are available with software drivers that interface with a host central processing unit to the graphics processor. In general, the software receives information for
5 drawing objects on a computer screen, calculates certain basic parameters associated with the objects and provides this to the graphics processor in the form of a "display list" of parameters. A graphics controller then uses the display list values in generating the graphics to be displayed. A graphics processor may
10 use interpolation techniques where the fundamental information for the object to be drawn comprises a series of initial and incremental parameters or values. The graphics processor loads or otherwise receives the initial parameters for the pixels to be drawn, interpolate the object by incrementing the parameters until
15 the object is completely drawn.

In many prior art computer systems, external devices such as graphics devices are able to read a stream of data (display list) from memory and execute programs stored in the memory in a similar manner. The size of these display list information tend to place
20 limitations on the traversal (read/write) speed of the central processing unit and the graphics processor.

The CPU typically builds the display list information with the instructions and parameters specific to the particular external device attached to the computer system. The external device then
25 reads the instruction stream and executes instructions from this stream. One of the common operations stored in the display list is

a command to load single and multiple registers of a device's register file with specified values.

Existing graphics implementation that use display lists typically load data in a sequential format to a register file in the graphics processor. For each type of primitive, a particular set of data values are required to render that type of primitive. For example, a point to be drawn to a pixel grid requires an X,Y location, color values and a Z value for depth comparison. An example of display list is shown below in Table I.

10
T,0040

<u>ADDRESS</u>	<u>NAME</u>	<u>DESCRIPTION</u>
0x4000	X	Initial X value
0x4004	Y	Initial Y value
0x4008	Z	Initial Z value
0x400C	R	Initial Red component
0x4010	G	Initial Green component
0x4014	B	Initial Blue component
0x4018	X1	Some other register
0x401C	X2	- -
0x4020	A	Alpha blending value

15
20

TABLE I

The display list in Table I provides the parameters required to draw points, lines and polygons. From the display list provided above, if a specific primitive rendering operation requires, for example, only the following register values to be loaded (e.g.,

X,Y,R,G,B and A); a prior art load instruction would use one of two alternative methods of instruction loading.

The first of the two alternatives will be to load all nine registers (e.g., "Load instruction (start at X),

5 X,Y,Z,R,G,B,X1,X2,A"). The stream of information in the display list will therefore occupy 10 instruction words (40 bytes) and load unnecessary registers.

The second load alternative is to use two consecutive load operations thereby replacing the two register load gaps (e.g.,
10 X1,X2) with only one load instruction (e.g., "Load instruction (start at X), X,Y,Z,R,G,B" and "Load instruction (starts at A), A"). The stream of information in the display list for this load sequence is 9 instruction words long (36 bytes). These two prior art instruction load methods have the common feature of

15 sequentially loading the register file with the parameter values for the primitive being rendered. Also, the load instructions comprise of two fields; a first field which holds the starting parameter value and a second field which holds the incremental count of subsequent parameter values for the primitive being
20 rendered.

Despite these prior methods, instruction load operations and the ability to load multiple registers contiguously to enable the efficient processing of the display list, several problems emerge when the size of the display list gets too large.

25 One of such problems is that extra system memory may be needed to store the large display list. This may impose extra cost in the

overall price of the computer system. Although memory prices are getting a bit cheaper, the average amount of memory installed in many of today's multimedia computer systems continue to substantially increase. For example, a Pentium® based multimedia computer system running MS Windows® NT may require about at least 32 megabytes of memory to run efficiently.

As the memory requirements of these multimedia systems continue to grow, the memory required to maintain and execute very long display lists needed by the multiprogramming operating systems in these computer systems become very significant. Moreover, since the memory in these systems may become locked, (i.e., the operating system is not able to swap processing to the computer system's external storage device). Such a lock further reduces the amount of memory that is left for the computer system to process other system activities.

Another problem with the presence of long display list is the time needed by the CPU to build the list and for the external device to execute the list. If a high frame rate and fast response time is needed by the CPU, the time spent managing the display list must be minimized. The amount of information that is being transferred between the CPU and the external device should not be sacrificed since the approach would definitely affect the quality of the image being rendered. Even if the setting ^{is} other than ^{for} ~~the~~ computer graphics, the amount of information may ~~have to~~ be the same since the external device may need it all.

As more and more of the computer's processing power is transferred to the central processing unit, the processing of long display lists to generate graphics display end up being bottlenecks in processing instructions by the CPU. This problem becomes even more pronounce if the processing of graphics data is transferred from a separate graphics processing chip ^{off} ~~of~~ device to the CPU.

Thus, a method of shortening display list information without losing the quality of the information being passed, while maintaining the processing speed of CPU is needed. The present invention provides the advantageous functionality of shortening display list information and the ability to randomly load register file in graphics processing device with a single load instruction.

Summary of the Invention

A method and apparatus are described herein which reduce processing time while maintaining the quality of display information and without requiring extra system memory. In accordance with the present invention, a graphics processor for generating shorter display list instructions without losing the quality of the display information supplied to a display screen is disclosed. The graphics processor provides a field load instruction which is generated by a central processing unit which is supplied to the graphics processor. ^{an} The field load instruction is then encoded into the display list instruction for subsequent execution by ~~a~~ external graphics device in a computer system. By providing a short display list, the present invention provides a system which is able to handle the increasing amount of graphics data processed in many present day multimedia computer systems, without requiring excessive amount of memory resources.

Another embodiment is a computer controlled graphics display system having a processor coupled to a bus, a memory unit coupled to the bus for storing the display list, a graphics processor for receiving microinstructions from the display list stored in the memory unit, a set of register files coupled to the graphics processor for storing the shortened display list in the graphics processor, and a private memory area disposed within the memory unit for storing address offsets of the display list; wherein named instructions generated by the central processor replace other means of randomly loading the register file in the graphics processor.

Embodiments further include the above; wherein the display list comprises parameterization procedures for processing polygon primitives, sets of graphics lines, and sets of graphics points; and wherein the parameterization procedure are further for
5 processing translation between different graphics formats.

Embodiments further include the above; wherein the load instruction comprises instruction bit-field for performing specific instructions by the display list.

Embodiment further include the above; wherein the load
10 instruction further comprises an opcode bit-field for storing data representing opcode instruction in the display list.

Embodiments further include the above; wherein the load instruction further comprise an partition bit-field for storing partition data defining the partition index of the display list to
15 the private memory area.

The graphics processor also preferably includes an internal instruction execution unit that receives the opcode from a prefetch unit and decodes the opcode. The execution unit also receives the display list and stores the display list in a register file.

Brief Description of the Drawings

Figure 1 is a simplified block diagram of a graphics processor coupled to a system bus of a computer system, in accordance with the principles of the present invention.

5 Figure 2 is a simplified block diagram showing in more detail a portion of the graphics subsystem of Figure 1.

Figure 3A is a simplified block diagram of the field load instruction unit Figure 2.

10 Figure 3B is a simplified block diagram of the partition look-up table of Figure 2.

Figure 4 is a flow diagram of the display list shortening process of the present invention.

08816053-031497

DETAIL DESCRIPTION OF THE PRESENT INVENTION

A method and apparatus for providing shorter display lists without losing the quality of the display information supplied to the graphics device is disclosed.

5 In the following detailed description of the present invention, numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be obvious to one skilled in the art that the present invention may be practiced without these specific details
10 or by using alternate elements or methods. In other instances well know methods, procedures, components, and circuits have been described in detail as not to unnecessarily obscure aspects of the present invention.

Some portions of the detailed description which follow are
15 represented in terms of procedures, logic blocks, processing, and other symbolic representations of operations on data bits within a computer system. These descriptions and representations are the means used by those skilled in the art to most effectively convey the substance of their work to other skilled in the art. A
20 procedure, logic block, process etc., is herein, and generally, conceived to be a self-consistent sequence of steps or instructions leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, magnetic signals capable of being stored, transferred, combined, compared,
25 and otherwise manipulated in a computer system. For reasons of convenience, and with reference to common usage, these signals are

referred to as bits, values or the like with reference to the present invention.

It should be borne in mind, however, that all of these terms are to be interpreted as referencing physical manipulations and quantities and merely convenient labels and are to be interpreted further in view of terms commonly used in the art. Unless specifically stated otherwise as apparent from the following discussions, it is understood that thorough discussions of the present invention, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data. The data is represented as physical (electronic) quantities within the computer system's registers and memories and is transformed into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

With reference to Figure 1, a block diagram is shown of a host computer system 100 used by the preferred embodiment of the present invention. In general, host computer 100 comprises a bus 101 for communicating data and instructions, a host processor (CPU) 102 coupled to bus 101 for processing data and instructions, a computer readable non-volatile memory unit 103 coupled to bus 101 for storing data and instructions from the host ^{a system memory 104} ~~processor 102~~, a computer readable data storage device 10⁵~~4~~ coupled to bus 101 for storing data and display device 106 coupled to bus 101 for

displaying information to the computer user. The display device 106 utilized with the computer system 100 of the present invention can be a liquid crystal device, cathode ray tube, or other display device suitable for creating graphics images and alphanumeric characters recognizable to the computer user.

The host system 100 provides data and control signals via bus 101 to a graphics hardware subsystem 109. ^(Figure 2) The graphics hardware 109 includes a ^{graphics} display processor 110 which executes a series of display instructions found within a display list. The graphics display processor 110 supplies data and control signals to a frame buffer which refreshes the display device for rendering images on display device. Alternatively, the host processor 102 may write the display list to the graphics processor 110 in accordance with known techniques.

It should be understood that the particular embodiment shown in Figure 1 is only one of many possible implementations of a graphics system for use in a computer system. Figure 1 is simplified for purposes of clarity so that many components and control signals are omitted which are not necessary to understand the present invention.

In the preferred embodiment, the graphics processor 110 provides hardware support for 2D and 3D graphics, and for text and windowing operations of a computer system. The graphics processor 110 transfers digital data from the system memory 104 or host processor 102, and processes data for storage in the RDRAM 115 ultimately for display on the display unit 106.

In accordance with the preferred embodiment, the host processor 102 provides necessary parameter values in the form of a display list, which typically is stored in system memory 104 until required by graphics processor 110.

5 The host processor 102 and system memory 104 both preferably communicate with the graphics processor 110 via the system bus 101. The system bus 101 may comprise any one a plurality of different types of host or input/output (I/O) buses, including the industry standard architecture (ISA), the extended ISA (EISA), the
10 peripheral component interconnect (PCI) and any other standardized system bus of a computer system.

Still referring to Figure 1, the graphics processor 110 couples to the system bus 101. In accordance with the preferred embodiment, the graphics processor 110 preferably includes bus mastering capabilities, thus permitting graphics processor 110 to bus master the system bus 101. Graphics processor 110 also couples
15 to ^{the} ~~a~~ display unit ¹⁰⁶ and a RDRAM 115. In the preferred embodiment, the RDRAM comprises a bank of RDRAM buffers, where the digital data stored in the RDRAM comprises a rectangular array of picture
20 elements referred to as pixels or pixel values. Each pixel can be defined by an 8 bit value, for example, which specifies the intensity of a single color of a corresponding pixel on a screen of the display unit 106.

The graphics device 109 hosts an array of volatile memory unit
25 referred to as register file 112. The register file 112 holds working information of the graphics device. The register file also
3

stores information and commands needed for operation of the graphics device 109.

The display unit 106 may be any suitable type of display device, such as a cathode ray tube (CRT) for desktop, workstation or server applications, a liquid crystal display (LCD) or any other suitable display device for a personal computer.

The RDRAM frame buffer provides a performance improvement by permitting faster access to display list instructions and pixel data, compared to accessing data stored in the main memory 104 of the host computer system 100. The graphics processor 110 communicates to the RDRAM buffer 115 through address data and control lines, collectively referred to as a RBUS 118.

Referring now to Figure 2, the graphics subsystem 109 preferably includes a register file 112, a graphics processor 110 and a frame buffer 115. Generally the register files 112 comprises a plurality of registers for storing the display list information. The register address generator generates the address pertaining to a register being accessed for display list information to be displayed.

The graphics processor 110 comprises a fetch subsequent parameters unit 200, a load instruction unit 210, a "right to left" shifter unit 220, an address counter 230 and a partition look-up table unit 240. ^{an address generator 235}

The field load instruction unit 210 comprises a plurality data bit locations for storing load bit data for performing the display list load instruction in the graphics processor. A detailed

description of the field load instruction is given in Figure 3A below.

The fetch subsequent instruction parameter unit 200 is coupled to the register files 112, and operates to fetch subsequent display list instructions after a first instruction has been processed. The fetch subsequent parameters unit is activated by the assertion of request for next parameter lines 201 by the graphics processor 110. When the fetch subsequent parameter 200 detects that request for the next parameter lines 201 have been asserted, display list data is driven on data line 221 to the register file 112 for subsequent write operation to the CPU.

Field Load instruction unit 210 is coupled to shifter 220 to pass load instructions to the register file 112. Field load instruction unit 210 comprises a plurality of data bits of a specified value each of which defines an operation to be performed by the graphics processor 110 in processing the display list. The field load instruction unit 210 passes data to shifter 220 when write enable signal lines 211 are asserted.

The Write Enable signal lines 211 are assumed to be the topmost bit position in shifter 220. At each internal clock cycle of the graphics processor 110, the Write Enable signal 211 is propagated to the register file 112 and to the subsequent parameter fetch unit 200 to fetch subsequent graphics parameters.

If the Write Enable signal 211 is asserted (i.e., having a bit value of "1'), the register file 112 stores the data provided by the fetch subsequent parameter unit 200 in a register address

provided by the address generation unit via the address counter 230.

If the Write Enable signal 211 is reset (i.e., having a binary value of "0"), all writes to the register file 112 are disabled and the subsequent parameter fetch unit 200 fetches new parameters from the display list. The shifter 220 shifts its contents one (1) bit to the left following either a write enable or a write disable operation to the register file 112. Shifting bits in the shifter 220, in this manner, allows the next bit of a Write Enable operation to generate a write/skip signal to the register file 112. Consequently, the register files 112 is randomly loaded depending on whether the write enable data bit is set or not.

Address counter 230 is coupled to the register file 112 and the address generation unit 235 to incrementally load new request addresses to the register file 112. The address counter 230 continues to generate new addresses to the register file 112 until the field load instruction contained in a display list are completely executed.

Still referring to Figure 2, partition look-up table 240 comprises a plurality of preloaded addresses which offset into the register file 112. The partition look-up table 240 is loaded with new address after each display list has been completely processed by the graphics processor 110. The partition table is coupled to the field load instruction unit. Portions of the field load instruction unit 210 reference the contents of the partition table 240.

In the preferred embodiment of the present invention, the partition lookup table 240 comprises 64 entries each of which is addressed by a partition data bit in the load instructions.

In its basic implementation, the look-up table 240 contains the addresses of 64 registers which are evenly distributed across the 1024 register set of the register files 112. Thus, each field load instruction only needs 6 bits to specify the starting partition of the register file 112 to load the display list thereby shortening the display list. The field load instruction also allows the register files 112 to be randomly loaded.

Referring to Figure 3A is a simplified block diagram of a load instruction of the preferred embodiment. The load instruction shown in Figure 3A comprises an opcode field 300, a write enable field 310 and a partition field 320.

The field load instruction of the preferred embodiment can load all, and only the registers required by a display list. The instruction stream of an exemplary load instruction looks as follows: "Field Load (write enables: 110111001), (partition starts at X), X,Y,R,G,B,A". This data stream, unlike the prior art, is only 7 instructions word long (28bytes). The write enable field 310 contained in the load instruction, which read from left to right, allows writes (binary 1s) only on desired registers of the register files 112. The registers that are not to be set are skipped.

Still referring to Figure 3A, the opcode field 300 stores data of a distinctive bit pattern which recognizes the "field load" instruction from other instructions in the display list information

by the graphics processor 110. In the preferred embodiment of the present invention, the opcode is kept short to leave more space for the "write enable" and the "partition" field respectively.

The write enable field 310 stores data bits which may be set to enable or disable register write operations of the load instruction to the register files 112. In the present invention, the setting of the write enable bit-field allows the register files 112 to be randomly loaded with the display parameter values. For example, if the write enable bit-field in a particular load instruction is enabled, the corresponding register location in the register files 112 is loaded with the display parameters.

Alternatively, if the write enable bit-field 310 is disabled, the write to the register files 112 will be disabled and the circuit which fetches subsequent parameters will request a next parameter fetch from the display list. Consequently, the corresponding register position is skipped in the register files 112. Thus, depending on the contents of the write enable bit-field position in a load instruction, corresponding register locations may be written or skipped.

The partition bit-field portion of the load instruction stores data bits which indexes to the partition look-up table.

Figure 3B is a simplified block diagram illustrating an exemplary embodiment of the partition look-up table of the present invention. The partition look-up table 240 shown in Figure 3B comprises of 64 entries of preloaded address offsets to the

register files 112. In the preferred embodiment, register files 112 comprises 1024 entries.

In order to address a particular register in the register files 112, prior art methods of addressing needed 10 binary bits of data to load each register. In the present implementation of the load instruction, partition look-up table 240 allows the register files 112 to be addressed with only 6 bits of data. The 64 entries in the partition look-up table 240 are evenly distributed across the register files 112 as shown in Figure 3B.

Referring to Figure 4, is a simplified block a flow process of the preferred embodiment of the present invention. The diagram shown in Figure 4 illustrates the execution of the "Field load" instruction. First at step 410, shifter 220 is loaded with the Write Enable data from the Write Enable field of the load instruction.

At step 420, the partition table is indexed using the partition instruction data bit from the load instruction. A base address of the first register in the register file 112 is then retrieved from the partition table 240 and loaded into the address counter 230 at step 430.

At step 440, the top bit of the shifter 220 is examined to determine whether the addressed register must be loaded or not. If the top-bit in shifter 220 is set, then the subsequent parameter fetch unit fetches the next parameter from the display list at step 450 and stores the retrieved data in the register file 112.

If the top-bit in shifter 220 is not set, then the address counter 230 increments the address count at step 460.

At step 470, shift register 220 is shifted one bit to the left after address counter 230 has been increased 1 bit. After the shifter has been shifted a bit, the contents of shifter 220 are examined to determine if it is empty at step 480. If the shifter 220 is empty processing of a current display list end at step 490. If, on the other hand, there is more data in the shifter 220, the graphics processor 110 continues to execute the current display list at step 440. The effect of loading the shifter 220 and incrementing the address count is to effectively load multiple register in the register file 112 randomly at once. Random loading of the register files 112 in this manner effectively shortens the display list compared to the traditional way of loading instructions in a sequential or contiguous manner.

Thus a method and an apparatus for shortening display list instruction through a random loading of register files is disclosed. The preferred embodiment of the present invention is described for illustrative purposes, numerous other variations of the disclosed embodiments will be apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such modification and variations.